arXiv:1406.6618v1 [stat.ML] 25 Jun 2014

When is it Better to Compare than to Score?

Nihar B. Shah **Sivaraman Balakrishnan** **Joseph Bradley**
Abhay Parekh **Kannan Ramchandran** **Martin Wainwright**
University of California, Berkeley

Abstract

When eliciting judgements from humans for an unknown quantity, one often has the choice of making direct-scoring (*cardinal*) or comparative (*ordinal*) measurements. In this paper we study the relative merits of either choice, providing empirical and theoretical guidelines for the selection of a measurement scheme. We provide empirical evidence based on experiments on Amazon Mechanical Turk that in a variety of tasks, (pairwise-comparative) ordinal measurements have lower per sample noise and are typically faster to elicit than cardinal ones. Ordinal measurements however typically provide less information. We then consider the popular Thurstone and Bradley-Terry-Luce (BTL) models for ordinal measurements and characterize the minimax error rates for estimating the unknown quantity. We compare these minimax error rates to those under cardinal measurement models and quantify for what noise levels ordinal measurements are better. Finally, we revisit the data collected from our experiments and show that fitting these models confirms this prediction: for tasks where the noise in ordinal measurements is sufficiently low, the ordinal approach results in smaller errors in the estimation.

1 Introduction

Eliciting judgements or knowledge about unknown quantities from non-expert humans is commonplace in many domains of society today. This has been facilitated by the emergence of several new 'crowdsourcing' platforms such as Amazon Mechanical Turk, that have become powerful, low-cost tools for collecting human knowledge and judgements. However, this low cost comes at the price of noise, due to the unreliability in the crowd response. This paper addresses this issue of noise at the source by studying how responses should be elicited.

We consider a setting in which humans perform evaluations that have numeric answers. Examples include a crowdsourcing task that involves counting the number of malaria parasites in an image of a blood smear [13], or a peer-grading task that involves students assigning grades to homeworks submitted by other students [18]. A standard design of such a task takes a *cardinal* approach where the evaluators directly enter numeric scores as the answers. This is illustrated by the example in Figure 1a where the subject is asked to rate the relevance of an image for the search query 'Internet' as a numeric entry between 0 and 100.

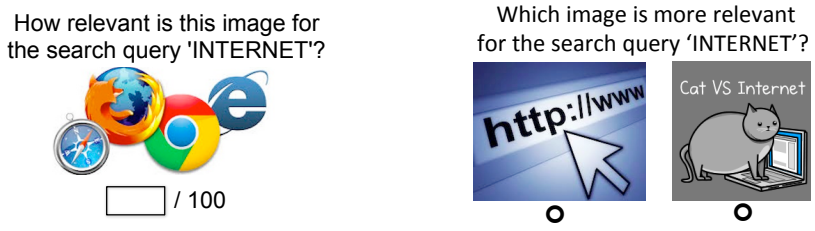
How relevant is this image for the search query 'INTERNET'?
/ 100
Which image is more relevant for the search query 'INTERNET'?
http://www
Cat VS Internet

(a) Cardinal interface

(b) Ordinal interface

Figure 1: Examples of cardinal and ordinal interfaces for a task on relevance rating.

Alternatively, one could take an *ordinal* approach, asking the evaluator to compare (or rank order) multiple items. Such an ordinal method is illustrated in Figure 1b where the evaluator is shown a pair of images, and is asked to select the one that is more relevant for the search query 'Internet'. In this paper, we restrict our attention to comparisons of only pairs of items in the ordinal setting.

Cardinal measurements allow for more precise measurements; in Figure 1, one cardinal measurement can take 100 values, whereas one ordinal measurement provides a single bit. One may be tempted to go even further and argue that ordinal measurements necessarily give less information, for one can always convert a set of cardinal measurements into ordinal, simply by ordering the measurements by value. The *data processing inequality* [5, Section 2.8] then suggests that an estimation procedure on any manipulation of the data cannot perform better than estimating from the original data. This may lead one to conclude that the ordinal data cannot yield superior results.

In contrast, ordinal measurements avoid calibration issues that are frequently encountered in cardinal measurements [23], such as the evaluators' inherent (and possibly time-varying) biases, or tendencies to give inflated or conservative evaluations. Ordinal measurements are also recognized to be easier or faster for humans to make [2, 20], allowing for more evaluations for the same level of time, effort and perhaps cost as well.

The lack of clarity regarding when to use a cardinal versus an ordinal approach forms the motivation for this paper. We first address the fundamental question of how much information we gain from each type of measurement. In extensive experiments on a variety of tasks, we find that the average per-sample noise is often significantly higher in cardinal measurements than in ordinal ones. In other words, *the data processing inequality does not apply when comparing cardinal and ordinal work from humans*.

While revealing, this still leaves two questions: Can we still make reliable estimates from paired comparisons? How much lower does the noise have to be for comparative measurements to be preferred over cardinal measurements? To address this, we invoke theoretical models for pairwise and cardinal measurements. We study the Thurstone (Case V) model [22], one of the most widely used models in both theory [4, 9] and practice [7, 19, 21]. We will show that it is indeed possible to perform estimation using pairwise comparisons, and via *minimax theory* we will quantify the settings in which pairwise comparisons are preferable to cardinal measurements. Minimax theory is a cornerstone of statistical decision theory and is a standard tool used in the comparison of estimators in a given model. In this paper, we will investigate the utility of this statistical perspective in comparing estimators *across* cardinal and ordinal models.

We also provide *topology-aware* bounds that incorporate the choice of pairs to be compared for the Thurstone and other popular pairwise-comparison models. Of particular importance is the popular Bradley-Terry-Luce (BTL) model [3, 12]. These bounds highlight the influence of the *comparison graph* on the estimation error.

Finally, we return to the data obtained from our experiments and fit our ordinal and cardinal models. We observe that the estimates produced from the ordinal data are more accurate than those from cardinal data when the ordinal noise is low enough. This suggests the following *practical guideline* in choosing between the cardinal and ordinal methods of data collection, of first estimating the noise in the two approaches by eliciting a few samples where the ground truth is known. The ordinal approach is then preferred if the ordinal noise is "low enough". For tasks in which the ordinal approach is preferred, our topology-aware results provide guidelines for the selection of items to compare when given a fixed budget.

2 Experiments Comparing Per-sample Noise in Cardinal and Ordinal

It is tempting to argue that a cardinal sample always gives more information than an ordinal sample: given cardinal samples, one can always order them thereby obtaining ordinal values. This argument suggests that an ordinal approach leads to a loss of information, and due to the data-processing inequality, cannot lead to better results. In this section, by means of seven different experiments conducted on Amazon Mechanical Turk (mturk.com), we show that such an argument is flawed. The experiments also provide insights into the per-sample noise in the ordinal and cardinal methods of data collection, which is a metric that the subsequent theory in this paper will also focus on.

Each experiment involved a certain task that was given to 100 human subjects. Each of these subjects was randomly given either the ordinal or the cardinal version of the task. Both versions had the same set of questions, and each question had a numeric answer. In the cardinal version of the task, the subject was

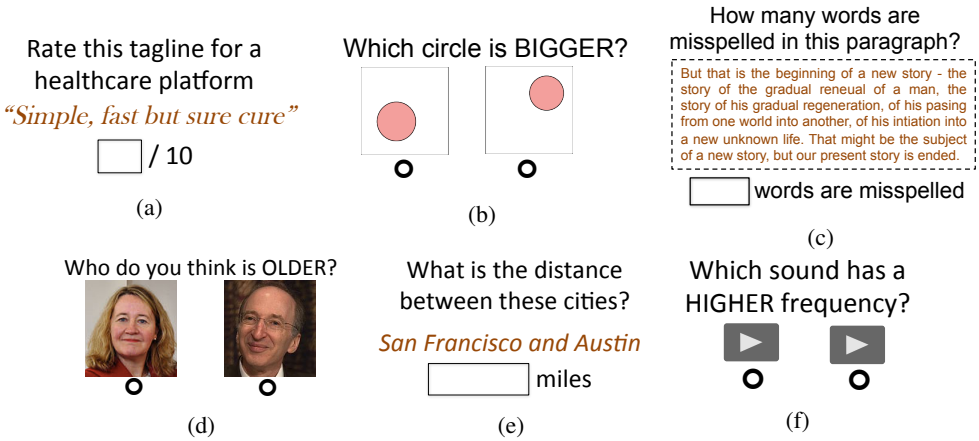
Rate this tagline for a healthcare platform
"Simple, fast but sure cure"
/ 10
(a)
Which circle is BIGGER?
(b)
How many words are misspelled in this paragraph?
But that is the beginning of a new story - the story of the gradual reneual of a man, the story of his gradual regeneration, of his pasing from one world into another, of his intiation into a new unknown life. That might be the subject of a new story, but our present story is ended.
words are misspelled
(c)
Who do you think is OLDER?
(d)
What is the distance between these cities?
San Francisco and Austin
miles
(e)
Which sound has a HIGHER frequency?
(f)

Figure 2: Screenshots of the tasks presented to the subjects. For each task, only one version (cardinal or ordinal) is shown here.

required to directly provide this number as the answer. The ordinal version presented the questions in pairs, and for each pair, the subject had to select the one which she believed had a larger number as the answer.

We now describe the tasks presented to the subjects in the seven experiments. The tasks were selected to have broad coverage of several important subjective judgment paradigms such as preference elicitation, knowledge elicitation, audio and visual perception, and skill utilization.

(a) Rating taglines for a product: A product was described and ten taglines for this product were shown (Figure 2a). The subject had to rate each of these taglines in terms of its originality, clarity and relevance to this product.

(b) Estimating areas of circles: The task comprised 25 questions. In each question, the subject was shown a circle in a bounding box (Figure 2b), and the subject was required to identify the fraction of the box's area that the circle occupied.

(c) Finding spelling mistakes in text: Eight paragraphs of text were shown, and the subject had to identify the number of words that were misspelled in each paragraph (Figure 2c).

(d) Estimating age of people from photographs: The subject was shown photographs of ten people (Figure 2d) and was asked to estimate the ages of the ten people.

(e) Estimating distances between pairs of cities: The subject was shown sixteen pairs of cities (Figure 2e) and for each pair, the subject had to estimate the distance between them.

(f) Identifying sounds: The subject was presented with ten audio clips, each of which was the sound of a single key on a piano (which corresponds to a single frequency). The subject had to estimate the frequency of the sound in each audio clip (Figure 2f).

(g) Rating relevance of the results of a search query: Twenty results for the query 'Internet' for an image search were shown (Figure 1) and the subject had to rate the relevance of these results with respect to the given query.

Upon obtaining the data from the experiments, we first reduced the cardinal data into ordinal form by comparing answers given by the subjects to consecutive questions. For five of the seven experiments ((b) through (f)), we had access to the "ground truth" solutions, using which we computed the fraction of answers that

Task	Tagline	Circle	Spelling	Age	Distance	Audio	Relevance
Error in Ordinal	**29%**	**6%**	**40%**	**13%**	**17%**	**20%**	**22%**
Error in Cardinal	31%	18%	46%	17%	46%	31%	27%
Time in Ordinal	**251s**	**98s**	**144s**	**31s**	**84s**	**66s**	**105s**
Time in Cardinal	342s	181s	525s	70s	305s	134s	185s

Table 1: Comparison of the average amount of error when ordinal data was collected directly vs. when cardinal data was collected and converted to ordinal. Also tabulated is the median time (in seconds) taken to complete a task by a subject in either type of task.

were incorrect in the ordinal and the cardinal-converted-to-ordinal data (any tie in the latter case was counted as half an error). For the two remaining experiments ((a) and (g)) for which there is no ground truth, we computed the 'error' as the fraction of (ordinal or cardinal-converted-to-ordinal) answers provided by the subjects that disagreed with each other.

The results are tabulated in Table 1 (boldface indicates a better performance). If the data-processing inequality were true, then it would be unlikely for the amount of error in the ordinal setting to be lower than that in the cardinal setting. On the contrary, one can see from Table 1 that converting cardinal data to an ordinal form results in a typically higher (and sometimes significantly higher) per-sample error than directly asking for ordinal evaluations. This absence of data-processing inequality may be explained by the argument that the inherent evaluation process in the human subjects is not the same in the cardinal and ordinal cases – humans do *not* perform an ordinal evaluation by first performing cardinal evaluations and then comparing them (this is why it is often found to be easier to compare than score [2, 20]). One can also see from Table 1 that the amount of time required for cardinal evaluations was typically (much) higher than for ordinal evaluations.

3 Theoretical Comparison of Cardinal and Ordinal Measurement Schemes

The experiments in the previous section established that the 'per-sample noise' in the cardinal setup is typically larger than that in the ordinal setting. However, each ordinal sample, unlike a cardinal value, can provide just one bit of information. This discrepancy is further complicated by the fact that the multitude of samples collected from multiple workers need to be aggregated in order to produce final estimates of the answers. It is thus not clear for a given a problem setting, whether an ordinal or a cardinal method of data collection would yield a superior performance. This section aims at addressing this issue: given that ordinal and cardinal samples have a different nature and amount of noise, which method of data collection will produce a smaller aggregate error?

In this section we focus our attention on the Thurstone (Case V) generative model [22], which is one of the most popular models considered in both theory [4, 9, 16] and practice [7, 19, 21]. This model assumes that every item has a certain numeric *quality score*, and a comparison of two items is generated via a comparison of the two qualities in the presence of an additive Gaussian noise.

We define a vector $\mathbf{w}^* \in \mathbb{R}^d$ of qualities, so item $j \in [d]$ has quality w_j^*. Under the Thurstone model we compare pairs of items items. For $i \in [n]$ the outcome of the i^{th} comparison is $y_i^{(o)} \in \{-1, 1\}$, where $y_i^{(o)}$ is given by

$$y_i^{(o)} = \text{sign}({\mathbf{w}^*}^T \mathbf{x}_i + \epsilon_i^{(o)}), \qquad \text{(THURSTONE)}$$

$\epsilon_i^{(o)}$ is independent Gaussian noise with variance σ_o^2, and $\mathbf{x}_i \in \mathbb{R}^d$ is a differencing vector with one entry $+1$, one entry -1 and the rest 0. Observe that the ordinal model is identifiable only upto a shift in $\mathbf{w}^*$ so we always assume $\mathbf{1}^T \mathbf{w}^* = 0$.

The cardinal analogue of this model involves a cardinal evaluation of individual items, where for $i \in [n]$ the outcome $y_i^{(c)}$ is given by

$$y_i^{(c)} = {\mathbf{w}^*}^T \mathbf{u}_i + \epsilon_i^{(c)} \qquad \text{(CARDINAL)}$$

where $\mathbf{u}_i$ in this case is a coordinate vector with one of its entries equal to 1 and remaining entries 0, and $\epsilon_i^{(c)}$ is independent Gaussian noise, with a *different variance* σ_c^2.

In order to build intuition on how to compare these models, in this section we focus on a simple scenario. Subsequently, in Section 4 we consider general settings. Analogous to the *fixed design regression* setup, we choose the vectors $\mathbf{x}_i$ a priori. Suppose that n is large enough, and that in the ordinal case we compare each pair $n/\binom{d}{2}$ times. In the cardinal case suppose that we evaluate the quality of each item n/d times.

To facilitate a comparison between the CARDINAL and THURSTONE models we consider the *minimax risk*. In each case a vector $\mathbf{w}$ induces a distribution $\mathbb{P}_{\mathbf{w}}$ from which the observed samples $\{y_1, \ldots, y_n\}$ are drawn (recall that the vectors $\mathbf{x}_i$ are *fixed*). Let $\mathcal{P}$ denote the family of induced distributions and $\mathcal{W}$ denote the set of allowed vectors $\mathbf{w}$. An estimator $\hat{\mathbf{w}}$ is a (measurable) map from the observed samples to $\mathcal{W}$. For a semi-norm ρ the minimax risk is

$$\mathfrak{M}_n^\rho := \inf_{\hat{\mathbf{w}}} \sup_{\mathbb{P}_{\mathbf{w}} \in \mathcal{P}} \mathbb{E}[\rho(\hat{\mathbf{w}}, \mathbf{w})]$$

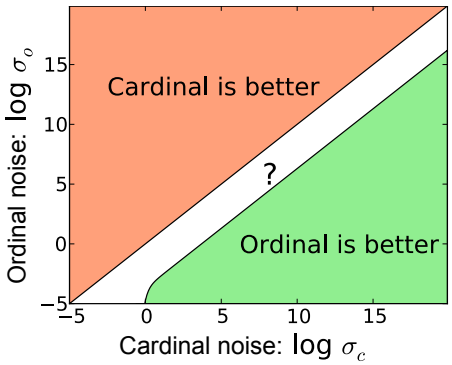

Figure 3: Characterizing the regions of (σ_c, σ_o) where cardinal or ordinal methods lead to a lower minimax error under the CARDINAL and THURSTONE models respectively. B is fixed at 1. The bounds for the THURSTONE model are loose when the signal to noise ratio (SNR) is high but relatively tighter at low SNR; the log-log scale of the axes attempts to focus on this low-SNR regime.

where the expectation is taken over the samples $\{y_1, \ldots, y_n\}$. The minimax risk characterizes the performance of the *best* estimator in the metric induced by ρ. In this section we focus on the case when $\rho(\hat{\mathbf{w}}, \mathbf{w}) = \|\hat{\mathbf{w}} - \mathbf{w}\|_2^2$ and we denote the minimax risk as $\mathfrak{M}_n^2$.

With these preliminaries in place we can attempt to ask a basic question for the simple case of evenly budgeted measurements: *Given n samples with noise standard-deviation σ_c in the cardinal case and σ_o in the ordinal case, is the expected minimax error in the estimation of d items lower in the cardinal case or the ordinal case?* The following theorem provides an answer for many regimes of (σ_c, σ_o).

Theorem 1 *Suppose that n is large enough and that in the* CARDINAL *model we observe each coordinate n/d times. The minimax risk is*

$$\frac{\mathfrak{M}_n^2(\text{CARDINAL})}{d} = \frac{d\sigma_c^2}{n}.$$

Suppose that n is large enough and that in the THURSTONE *model we observe each pair $n/\binom{d}{2}$ times. Suppose $\|\mathbf{w}^*\|_\infty \leq B$, and that B and σ_o are known. Let Φ denote the standard Gaussian c.d.f., and let $\kappa := \Phi(2B/\sigma_o)(1 - \Phi(2B/\sigma_o))$. Then the minimax risk is bounded as*

$$0.0008\kappa\frac{d\sigma_o^2}{n} \leq \frac{\mathfrak{M}_n^2(\text{THURSTONE})}{d} \leq \frac{5}{\kappa^2}\frac{d\sigma_o^2}{n}.$$

In the cardinal case when each coordinate is measured the same number of times, the CARDINAL model reduces to the well-studied normal location model, for which the MLE is known to be the minimax estimator and its risk is straightforward to characterize (see [10] for instance). In the ordinal case the result follows from the general treatment in Section 4. Observe that the THURSTONE minimax bounds depend on $\|\mathbf{w}^*\|_\infty$. This is related to the strong convexity parameter of the likelihood in the THURSTONE model which degrades for increasing $\|\mathbf{w}^*\|_\infty$. Informally, this is related to the difficulty of estimating very small (or very large) probabilities that can arise in the THURSTONE model for large $\|\mathbf{w}^*\|_\infty$.

Observe from Theorem 1 that the minimax risks in the cardinal and ordinal settings have the same dependency on d and n. An ordinal approach of collecting data is thus better overall whenever its per-sample error is "low enough". Figure 3 summarizes the result of Theorem 1.

4 General Bounds and Topology Considerations

In the previous section, we analyzed one paired comparison model, the THURSTONE model. We now provide a more general treatment by considering three models while allowing arbitrary comparisons. In addition to the THURSTONE model we also provide results for its linear and logistic analogues:

$$y_i = {\mathbf{w}^*}^T\mathbf{x}_i + \epsilon_i \qquad \text{for } i \in [n], \tag{PAIRED LINEAR}$$

where ϵ_i are i.i.d. $N(0, \sigma^2)$, and

$$\mathbb{P}(y_i = 1|\mathbf{x}_i, \mathbf{w}^*) = \frac{1}{1 + \exp\left(\frac{-{\mathbf{w}^*}^T\mathbf{x}_i}{\sigma}\right)} \qquad \text{for } i \in [n]. \tag{BTL}$$

As before, the $\mathbf{x}_i$'s are difference vectors, and we assume $\mathbf{1}^T\mathbf{w}^* = 0$. The second model is the popular Bradley-Terry-Luce (BTL) model [3, 12]. The BTL model is also a popular choice for modeling pairwise

comparisons [1, 6, 8, 11, 16], especially since it allows for a computationally simple maximum likelihood inference. The parameter σ plays the role of a noise parameter, with a higher value of σ leading to more uncertainty in the comparisons. We will assume, under all the models, that the value of σ is known. We note in passing that each of these models is a special case of *generalized linear models* (GLMs) [14], and that many of the insights here carry over to this general class. We defer a detailed treatment of GLMs to an extended version.

In this section we will not assume that items are chosen uniformly at random, rather we provide bounds in the general case when the measurements are fixed *a priori*. This will highlight the central role of the Laplacian of the weighted graph of chosen comparisons. The minimax rate for estimating the underlying quality in ℓ_2^2 will depend on the spectral properties of the Laplacian which in turn depends on the *topology* of the underlying comparison graph.

In the ordinal models, each measurement is related to a *difference* of two quality assessments. Observe that the *covariance matrix* of the measurements is

$$\hat{\Sigma} := \frac{1}{n}\sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^T := \frac{L}{n}$$

where L is the combinatorial graph Laplacian of the undirected graph with each edge having a weight equal to the number of times its end points are compared. We refer to $\hat{\Sigma}$ as the *standardized* Laplacian. The standardized Laplacian is positive semi-definite and has at least one zero-eigenvalue corresponding to the all ones vector. We assume that the graph induced by the comparisons is *connected*, since it is easy to verify that without this the model is not identifiable. The covariance matrix induces a semi-norm on vectors in $\mathbb{R}^d$, defined as $\|\mathbf{v}\|_{\hat{\Sigma}} = \sqrt{\mathbf{v}^T \hat{\Sigma} \mathbf{v}}$. We denote the Moore-Penrose pseudo inverse of $\hat{\Sigma}$ by $\hat{\Sigma}^\dagger$. We first focus on the minimax risk of estimating $\mathbf{w}^*$ in the *squared* semi-norm induced by $\hat{\Sigma}$. We denote this as $\mathfrak{M}_n^{\hat{\Sigma}}$. Theorem 2 below bounds this minimax risk in each of the three models. To cleanly state our results we make the simplifying assumption that $d > 9$.

Theorem 2.A (PAIRED LINEAR**)** *The minimax rate is bounded as*

$$0.00013\frac{d\sigma^2}{n} \le \mathfrak{M}_n^{\hat{\Sigma}}(\text{PAIRED LINEAR}) \le 0.68\frac{d\sigma^2}{n} .$$

Theorem 2.B (THURSTONE**)** *Assume that $\|\mathbf{w}^*\|_\infty \le B$ (known). Let $\kappa := \Phi(2B/\sigma)(1-\Phi(2B/\sigma))$, and let $n \ge \frac{\sigma^2 \kappa \mathrm{tr}(\hat{\Sigma}^\dagger)}{0.035 B^2}$. The minimax rate is bounded as*

$$0.0008\kappa\frac{d\sigma^2}{n} \le \mathfrak{M}_n^{\hat{\Sigma}}(\text{THURSTONE}) \le \frac{5}{\kappa^2}\frac{d\sigma^2}{n} .$$

Theorem 2.C (BTL**)** *Assume that $\|\mathbf{w}^*\|_\infty \le B$ (known) and $n \ge \frac{0.04467\sigma^2 \mathrm{tr}(\hat{\Sigma}^\dagger)}{B^2}$. The minimax rate is bounded as*

$$0.001\frac{d\sigma^2}{n} \le \mathfrak{M}_n^{\hat{\Sigma}}(\text{BTL}) \le 1.37\left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^4 \frac{d\sigma^2}{n} .$$

The upper bound in each case is from an analysis of the maximum likelihood (ML) estimator. The ML estimator, in all three settings, is the solution to a convex-optimization problem (while this is clear for the PAIRED LINEAR and BTL models, see for instance [23] for a proof in the THURSTONE case).

Proof Sketch: Lower bound: The lower bounds are based on a combination of information-theoretic techniques and carefully constructed packings of the parameter set $\mathcal{W}$. Such techniques are standard in minimax analysis [24]. The main technical difficulty is in constructing a packing in the semi-norm induced by $\hat{\Sigma}$. A consequence of Fano's inequality (see for instance Theorem 2.5 in [24]) is that if we can construct a packing of vectors $\{\mathbf{w}_1, \ldots, \mathbf{w}_M\}$ of vectors in $\mathcal{W}$ such that (a) the KL divergence between the induced distributions is small, i.e. $\max_{ij} D_{\mathrm{KL}}(\mathbb{P}_{\mathbf{w}_i} \| \mathbb{P}_{\mathbf{w}_j}) \le \beta \log M$ for a sufficient small (universal) constant β, and (b) $\min_{ij} \|\mathbf{w}_i - \mathbf{w}_j\|_{\hat{\Sigma}}^2 \ge \delta$ for some parameter δ, then for a small constant c, the minimax risk above is at least $c\delta$. The main effort is in constructing an *exponentially* large (in d) packing in the $\hat{\Sigma}$ norm with sufficiently large δ, and bounding the model specific constants β and c above. The condition on n is used to ensure that the constituents of the packing satisfy $\|\mathbf{w}\|_\infty \le B$. We relegate the details to the Appendix.

Upper bound: In each case we analyze the maximum likelihood estimator $\hat{\mathbf{w}} = \arg\min_{\mathbf{1}^T\mathbf{w}=0} \ell(\mathbf{w})$ where $\ell : \mathcal{W} \mapsto \mathbb{R}$ is the negative log-likelihood under the corresponding model. In the case of the BTL and THURSTONE models we impose the additional constraint that $\|\mathbf{w}\|_\infty \leq B$.

The optimization problem in each case is *convex*. The analysis follows along the lines of standard statistical analyses of M-estimators [25]. We proceed by upper and lower bounding the quantity

$$f(\hat{\mathbf{w}}) = \ell(\hat{\mathbf{w}}) - \ell(\mathbf{w}^*) - \langle \nabla\ell(\mathbf{w}^*),\, \hat{\mathbf{w}} - \mathbf{w}^* \rangle$$

where $\nabla\ell(\mathbf{w}^*) \in \mathbb{R}^d$ is the gradient of the negative log-likelihood.

In particular, an analysis of the *strong convexity* parameter of the negative log-likelihood provides a lower bound of the form $f(\hat{\mathbf{w}}) \geq \gamma\|\hat{\mathbf{w}} - \mathbf{w}^*\|^2_{\hat{\Sigma}}$ for an appropriate γ. Since $\hat{\mathbf{w}}$ is the maximum-likelihood estimator, we get $\ell(\hat{\mathbf{w}}) \leq \ell(\mathbf{w}^*)$. This implies $f(\hat{\mathbf{w}}) \leq -\langle \nabla\ell(\mathbf{w}^*),\, \hat{\mathbf{w}} - \mathbf{w}^* \rangle \leq \|\hat{\mathbf{w}} - \mathbf{w}^*\|_{\hat{\Sigma}} \|\nabla\ell(\mathbf{w}^*)\|_{\hat{\Sigma}^\dagger}$ via Cauchy-Schwarz under appropriate conditions (recall that Σ only induces a semi-norm). Putting these together we arrive at the bound, $\gamma\|\hat{\mathbf{w}} - \mathbf{w}^*\|_{\hat{\Sigma}} \leq \|\nabla\ell(\mathbf{w}^*)\|_{\hat{\Sigma}^\dagger}$. The main model-specific effort is in analyzing the strong convexity parameter and bounding the $\Sigma^\dagger$-norm of $\nabla\ell(\mathbf{w}^*)$. We defer the details to the Appendix. ■

A minimax analysis of the BTL model is also provided in Negahban et al. [15]. Although their main focus is the analysis of a random walk based algorithm, they also provide an analysis for the MLE for the case of uniformly randomly chosen $\mathbf{x}_i$. Their information theoretic lower bound studies a related but different problem. Their analysis applies only to the specific sampling schemes considered and show a considerable gap between the MLE and the lower bound. Our analysis however eliminates this discrepancy and shows that MLE is in fact minimax (rate) optimal for $\mathfrak{M}_n^{\hat{\Sigma}}$.

To conclude this section, let us develop some consequences of this theorem. Let us focus on upper bounds in the ordinal setting, and consider estimation error in ℓ_2^2. As in the theorem, we assume that the graph induced by the comparisons is connected. Now ignoring model specific constants we can see that

$$\mathfrak{M}_n^2 \leq \frac{d\sigma^2}{n\lambda_2(\hat{\Sigma})}$$

where $\lambda_2(\hat{\Sigma})$ is the second smallest eigenvalue of $\hat{\Sigma}$. Recall that $\hat{\Sigma}$ is simply the standardized Laplacian of the comparison graph, and its second eigenvalue is determined by the *topology* of the chosen comparisons. To understand this we consider three canonical examples, and in each case we assume that the comparison graph is fixed, n is large enough and that the samples are distributed evenly along the fixed graph. It is straightforward to extend this to the case of *randomly* chosen comparisons from a fixed graph using matrix concentration inequalities (see for instance [17]).

1. Dumbell graph: This is the graph on d vertices, which consists of two cliques of $d/2$ disjoint sets of vertices with a single edge between them. Suppose $n \geq \binom{d/2}{2}+1$. Since the unweighted graph has $\lambda_2 = O(1)$ we get $\lambda_2(\hat{\Sigma}) = \frac{O(1)}{\binom{d/2}{2}+1}$ and the ℓ_2^2 error scales as $\mathfrak{M}_n^2 \leq \frac{\binom{d/2}{2} d\sigma^2}{n}$.

2. Complete graph: Suppose $n \geq \binom{d}{2}$. It is easy to verify that since the unweighted complete graph has $\lambda_2 = d$, we get $\lambda_2(\hat{\Sigma}) = \frac{d}{\binom{d}{2}}$ and the ℓ_2^2 error scales as $\mathfrak{M}_n^2 \leq \frac{\binom{d}{2}\sigma^2}{n}$.

3. Degree-k expander: The unweighted degree-k expander has $\lambda_2 = O(k)$ and a similar argument as before shows that if $n \geq kd$ then we get the error scales as $\mathfrak{M}_n^2 \leq \frac{d^2\sigma^2}{n}$.

To summarize we see the ℓ_2^2 error scaling of $\frac{d^2\sigma^2}{n}$ for the complete graph and the degree-k expander. We conjecture that this is in fact the best possible scaling. Observe that the degree-k expander requires $n \geq kd$ while the complete graph requires $n \geq \binom{d}{2}$, so in practical applications at least for small sample sizes, we should prefer a low-degree expander. On the other hand, for the dumbell graph, the error scales as $d^3\sigma^2/n$ indicating that is a bad topology.

5 Inference in the Experimental Data

In this section we return to our experimental data from Section 2. We consider data from the three experiments of identifying number of spelling errors, estimating the distances between cities, and recognizing the

frequencies of audio, for which we know the ground truth. For each of the three experiments, we execute 100 iterations of the following procedure. Select five workers from the cardinal and five from the ordinal pool of workers who did this experiment, uniformly at random without replacement. (The number five is inspired by practical systems [18, 26].) Run the maximum-likelihood estimator of the CARDINAL model on the data from the five workers selected from the cardinal pool, and the maximum-likelihood estimator of the THURSTONE model on the data from the five workers of the ordinal pool. In particular, the estimator for the ordinal case first estimates σ via 3-fold cross-validation, choosing the value that maximizes held-out data log likelihood, and then uses this best fit for the rest of the estimation procedure. Note that unlike Section 2, the cardinal data here is *not* converted to ordinal.

We evaluated the performance of these two estimators as follows. The true and inferred vectors were first scaled to have their maximum elements equal to 1 and minimum elements equal to -1; this mimics the effect of knowing the scaling B via 'domain knowledge'. The (scaled) inferred vectors in either case were then compared with the (scaled) true vector in terms of two metrics: (i) $\frac{1}{d}$ times the squared ℓ_2 distance, and (ii) the Kendall's tau rank correlation coefficient.

The results of this evaluation are enumerated in Table 2 (boldface indicates a better performance). To put the results in perspective of the rest of the paper, let us also recall the per-sample errors in these experiments from Table 1. Observe that in the experiment of estimating distances, the per-sample error in the cardinal data was significantly higher than the ordinal data. This is reflected in the results of Table 2 where the estimator on the ordinal data performs much better (in terms of the ℓ_2 error) than the estimator on the cardinal data. On the other hand, the task of identifying the number of spelling mistakes involved a per-sample noise that was comparable across the two settings, and hence the estimator on the cardinal data scores over the ordinal one. As one would expect, the ordinal approach outperforms cardinal in terms of the (ordinal) Kendall's tau coefficient.

Task	Spelling	Distance	Audio
Squared ℓ_2-distance in Ordinal	0.358	**0.168**	**0.444**
Squared ℓ_2-distance in Cardinal	**0.350**	0.330	0.508
Kendall's tau coefficient in Ordinal	**0.277**	**0.547**	**0.513**
Kendall's tau coefficient in Cardinal	0.129	0.085	0.304

Table 2: Evaluation of the inferred solution from the data received from multiple workers.

6 Conclusion

This paper compares cardinal and ordinal approaches to evaluation performed by humans. With an increasing number of systems relying on non-expert human evaluators (e.g., using crowdsourcing), the choice of the evaluation mechanism forms a critical component of these systems. We argue by means of experiments and fundamental theoretical bounds that ordinal data provides a better estimate of the true solution when the per-sample noise is low enough relative to cardinal data, and the threshold for this choice is independent of the number of observations and the number of questions. This suggests a guideline for deciding whether to deploy a cardinal or an ordinal method of data collection: estimate the noise in the data by obtaining a few samples from either method, and then use the bounds on the overall error to determine the better of the two options.

We suggest further research to understand the tradeoffs in cardinal and ordinal measurements. Our theoretical results were based on simple models, but more complex models, such as ones incorporating the abilities of the different human workers, could be more accurate. Other model classes might have different noise thresholds determining when cardinal or ordinal performs best. Also, it would be useful to make in-depth studies of noise in specific crowdsourcing settings, such as user experience testing and peer grading in classes.

Future research could also improve data collection. For both cardinal and ordinal data, it would be useful to derive methods for adaptively choosing which measurements to take. Our results on topology-aware bounds could potentially be used to improve ordinal evaluation by analyzing the best topologies for choosing pairs of items to compare.

References

[1] D. R. Atkinson, B. E. Wampold, S. M. Lowe, L. Matthews, and H.-N. Ahn. Asian American preferences for counselor characteristics: Application of the Bradley-Terry-Luce model to paired comparison data. *The Counseling Psychologist*, 26(1):101–123, 1998.

[2] W. Barnett. The modern theory of consumer behavior: Ordinal or cardinal? *The Quarterly Journal of Austrian Economics*, 6(1):41–65, 2003.

[3] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pages 324–345, 1952.

[4] T. Bramley et al. A rank-ordering method for equating tests by expert judgment. *Journal of Applied Measurement*, 6(2):202–223, 2005.

[5] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2012.

[6] S. Heldsinger and S. Humphry. Using the method of pairwise comparison to obtain reliable teacher assessments. *The Australian Educational Researcher*, 37(2):1–19, 2010.

[7] R. Herbrich, T. Minka, and T. Graepel. Trueskill: A bayesian skill rating system. *Advances in Neural Information Processing Systems*, 19:569, 2007.

[8] K. J. Koehler and H. Ridpath. An application of a biased version of the Bradley-Terry-Luce model to professional basketball results. *Journal of Mathematical Psychology*, 25(3), 1982.

[9] P. F. Krabbe. Thurstone scaling as a measurement method to quantify subjective health outcomes. *Medical care*, 46(4):357–365, 2008.

[10] E. Lehmann and G. Casella. *Theory of Point Estimation*. Springer Texts in Statistics. 1998.

[11] P. J. Loewen, D. Rubenson, and A. Spirling. Testing the power of arguments in referendums: A Bradley–Terry approach. *Electoral Studies*, 31(1):212–221, 2012.

[12] R. D. Luce. Individual choice behavior, a theoretical analysis. *Bull. Amer. Math. Soc. 66 (1960), 259-260*, pages 0002–9904, 1960.

[13] M. A. Luengo-Oroz, A. Arranz, and J. Frean. Crowdsourcing malaria parasite quantification: an online game for analyzing images of infected thick blood smears. *Journal of medical Internet research*, 14(6), 2012.

[14] P. McCullagh and J. Nelder. *Generalized Linear Models*. Monographs on Statistics and Applied Probability. Routledge, Chapman & Hall, Incorporated, 1983.

[15] S. Negahban, S. Oh, and D. Shah. Rank centrality: Ranking from pair-wise comparisons. *arXiv preprint arXiv:1209.1688*, 2014.

[16] R. M. Nosofsky. Luce's choice model and Thurstone's categorical judgment model compared: Kornbrot's data revisited. *Attention, Perception, & Psychophysics*, 37(1):89–91, 1985.

[17] R. I. Oliveira. Concentration of the adjacency matrix and of the laplacian in random graphs with independent edges, 2009.

[18] C. Piech, J. Huang, Z. Chen, C. Do, A. Ng, and D. Koller. Tuned models of peer assessment in MOOCs. In *International Conference on Educational Data Mining*, 2013.

[19] D. Ross. Arpad Elo and the Elo rating system, 2007.

[20] N. Stewart, G. D. Brown, and N. Chater. Absolute identification by relative judgment. *Psychological review*, 112(4):881, 2005.

[21] J. Swets. The relative operating characteristic in psychology. *Science*, 182(4116), 1973.

[22] L. L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273, 1927.

[23] K. Tsukida and M. R. Gupta. How to analyze paired comparison data. Technical report, DTIC Document, 2011.

[24] A. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. 2008.

[25] S. van de Geer. *Empirical Processes in M-Estimation*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2009.

[26] J. Wang, P. G. Ipeirotis, and F. Provost. Managing crowdsourcing workers. In *The 2011 Winter Conference on Business Intelligence*, pages 10–12, 2011.

Appendix A provides some additional details on the experiments. Appendix B reviews some technical results that are used in our theoretical proofs. Appendix C presents proofs of the theoretical results.

A Additional Details on Experiments

This section presents additional details on the experiments presented in Section 2 and Section 5. We first discuss the experiments of Section 2. The data was collected by putting up tasks on Amazon Mechanical Turk (mturk.com). Amazon Mechanical Turk is an online platform for putting up tasks, where any individual or institution can put up tasks and offer certain payments, and anyone can log in and complete the tasks in exchange for some payment that was specified along with the task. The following are some additional specifics about the experiments described in Section 2.

- Each experiment comprised of 100 tasks, all comprising the same set of questions but organized in either a cardinal or ordinal format at random.
- A worker was offered 20 cents for any task she completed.
- A worker was allowed to do no more than one task in an experiment.
- Workers were required to answer all the questions in a task.
- Only those workers who had 100 or more approved works prior to this and also had at least 95% approval rate were allowed.
- Workers from any country were allowed to participate, except for the task of estimating distances between cities where only workers from the USA were allowed since all the questions were about American cities.

We now move on to discuss the inference algorithms of Section 5. The inference algorithms in this section operated on the data from three of the experiments. The four remaining experiments were unsuitable for this purpose: the experiments on rating the relevance of search results and rating taglines for a product had no ground truth; the task of identifying the area of circles had each question drawn independently at random from a Beta distribution, and hence no two workers answered the same questions; the comparison graph for the experiment on identifying age from pictures was not a connected graph.

Table 2 in Section 5 presented the average errors across 100 runs of the inference procedure; Table 3 here tabulates the associated standard deviation of the errors across the 100 runs.

Task	Spelling	Distance	Audio
Squared ℓ_2-distance in Ordinal	0.122	0.070	0.302
Squared ℓ_2-distance in Cardinal	0.207	0.076	0.279
Kendall's tau coefficient in Ordinal	0.244	0.113	0.217
Kendall's tau coefficient in Cardinal	0.214	0.148	0.239

Table 3: The standard deviation of the errors incurred in the 100 runs of the inference procedure of Section 5. The average of the errors is listed in Table 2.

B Review of some Technical Results

In this section we present some well known information-theoretic results that we use in our proofs. See for instance [3] for proofs of these claims.

B.1 Fano's inequality: Multiple hypothesis version

Lemma 3 *Let X be a random variable with distribution equal to one of $r+1$ possible distributions $\mathbb{P}_1, \ldots, \mathbb{P}_{r+1}$. Furthermore, the Kullback-Leibler divergence between any pair of densities cannot be too large,*

$$D_{\mathrm{KL}}(\mathbb{P}_i \| \mathbb{P}_j) \leq \beta \ \forall \ i \neq j.$$

Let $\psi(X) \in \{1, \ldots, r+1\}$ be an estimate of the index. Then

$$\sup_i \mathbb{P}_i(\psi(X) \neq i) \geq 1 - \frac{\beta + \log 2}{\log r}.$$

B.2 Fano's inequality: Two hypothesis version

Lemma 4 *Let X be a random variable with distribution either $\mathbb{P}_0$ or $\mathbb{P}_1$, and suppose that the KL divergence between f_0 and f_1 be bounded as*

$$D_{\mathrm{KL}}(\mathbb{P}_0 \| \mathbb{P}_1) \leq \alpha < \infty.$$

Then for any $\psi(X) \in \{0,1\}$ we have

$$\sup_{i \in \{0,1\}} \mathbb{P}_i(\psi(X) \neq i) \geq \frac{1 - \sqrt{\alpha/2}}{2}.$$

B.3 Estimation error

Let $\mathcal{P}$ be a family of distributions. Consider a map $\theta : \mathcal{P} \mapsto \Omega$, and let ρ be a semi-norm on Ω. An estimator $\hat{\theta}$ is a measurable function $\hat{\theta} : \mathcal{P} \mapsto \Omega$. The following Lemma gives a lower bound on the minimax error in estimating θ in the metric induced by ρ.

Lemma 5 *Let X be a random variable with distribution equal to one of $r+1$ possible distributions $\mathbb{P}_1, \ldots, \mathbb{P}_{r+1}$ such that*

$$D_{\mathrm{KL}}(\mathbb{P}_i \| \mathbb{P}_j) \leq \beta \ \ \forall \ \ i \neq j.$$

Suppose that

$$\min_{ij} \rho(\theta(\mathbb{P}_i), \theta(\mathbb{P}_j)) \geq \delta$$

then the minimax estimation error of any estimator is lower bounded as

$$\inf_{\hat{\theta}} \sup_{i \in \{1,\ldots,r+1\}} \mathbb{E}[\rho(\hat{\theta}, \theta(\mathbb{P}_i))] \geq \frac{\delta}{2}\left(1 - \frac{\beta + \log 2}{\log r}\right).$$

C Proofs

We first introduce some notation which will be employed subsequently in the proofs. Observe that L is a positive semi-definite matrix (recall from Section 4). Let

$$L = U \Lambda U^T \, .$$

Let $\lambda_1 > \ldots > \lambda_d$ be the eigenvalues of L and assume without loss of generality that $\forall i \in [d]$, λ_i is the $(i,i)^{\text{th}}$ entry of Λ. Since the graph topologies are assumed to be connected, we have $\lambda_i \neq 0 \, \forall \, i \in [d-1]$ and $\lambda_d = 0$. The Moore-Penrose pseudoinverse of L is the $(d \times d)$ matrix K, and this satisfies

$$K := U \tilde{\Lambda} U^T \quad \text{where } \tilde{\lambda}_i = \lambda_i^{-1} \mathbf{1}\{\lambda_i \neq 0\}.$$

Note that K is also positive semidefinite, has a rank equal to $(d-1)$, and $Tr(LK) = d-1$. Furthermore, $\Lambda \tilde{\Lambda} = \tilde{\Lambda}^{\frac{1}{2}} \Lambda \tilde{\Lambda}^{\frac{1}{2}} = \sum_{i=1}^{d-1} \mathbf{e}_i \mathbf{e}_i^T$.

The standardized versions of L and K are $\hat{\Sigma} := \frac{1}{n} L$ and $\hat{\Sigma}^\dagger := nK$. Note that $\hat{\Sigma}^\dagger$ is the Moore-Penrose pseudoinverse of $\hat{\Sigma}$.

We first state two lemmas that we will use to prove our results. Lemma 6 is used to prove lower bounds and Lemma 7 is used to prove upper bounds. The proofs of the two Lemmas are provided at the end of this section.

Lemma 6 *For any $\delta > 0$, $\alpha \in (0,1)$, $\beta \in (0,1)$ with $\beta = \frac{\log 2 + \alpha \log \alpha - \alpha}{2}$, there exist a set of $e^{\beta d}$ vectors $\{\mathbf{w}_1, \ldots, \mathbf{w}_{e^{\beta d}}\}$, each of length d, such that every pair of vectors satisfies*

$$\alpha \delta^2 \leq (\mathbf{w}_i - \mathbf{w}_j)^T M (\mathbf{w}_i - \mathbf{w}_j) \leq 4\delta^2$$

and every vector in this set also satisfies

$$\mathbf{1}^T \mathbf{w}_i = 0 \, .$$

Lemma 7 *Consider any $(d \times d)$ positive semidefinite matrix L, and any vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$ such that $\mathbf{x} \perp$ nullspace(L). If K is the Moore-Penrose pseudoinverse of L, then*

$$\mathbf{x}^T\mathbf{y} \leq \sqrt{\mathbf{x}^T L \mathbf{x}}\sqrt{\mathbf{y}^T K \mathbf{y}} \, .$$

Proof of Theorem 1 In the cardinal case when each coordinate is measured the same number of times, the CARDINAL model reduces to the well-studied normal location model, for which the MLE is known to be the minimax estimator and its risk is straightforward to characterize (see [2] for instance).

In the ordinal case the result follows from Theorem 2.B, with $\hat{\Sigma} = \frac{2}{d(d-1)}\left(dI - \mathbf{1}\mathbf{1}^T\right)$, i.e., an appropriately scaled Laplacian of the complete graph. We know that $\mathbf{1}^T\mathbf{w}^* = 0$, and further observe in the proof of Theorem 2.B that it suffices to consider $\hat{\mathbf{w}}$ such that $\mathbf{1}^T\hat{\mathbf{w}} = 0$. It follows that $\mathfrak{M}_n^{\hat{\Sigma}}(\text{THURSTONE}) = \frac{2d}{d-1}\frac{\mathfrak{M}_n^2(\text{THURSTONE})}{d}$. The quantity $\mathfrak{M}_n^{\hat{\Sigma}}(\text{THURSTONE})$ is bounded in Theorem 2.B. ■

Proof of Theorem 2.A (PAIRED LINEAR):
Lower Bounds: For any $\mathbf{w}_1$ and $\mathbf{w}_2$, the KL divergence between the distributions of $\mathbf{y}$ under $\mathbf{w}_1$ and $\mathbf{w}_2$ as the true values is

$$D_{\mathrm{KL}}(P_{\mathbf{w}_1}(\mathbf{y}) \| P_{\mathbf{w}_2}(\mathbf{y})) = \frac{1}{\sigma^2}(\mathbf{w}_1 - \mathbf{w}_2)^T L (\mathbf{w}_1 - \mathbf{w}_2) \, .$$

For any $\delta > 0$, Lemma 6 constructs a packing $\{\mathbf{w}_1, \ldots, \mathbf{w}_{e^{\beta d}}\}$ such that every pair of distinct vectors $\mathbf{w}_i$ and $\mathbf{w}_j$ in this packing satisfies (with $\alpha = 0.15$ and $\beta = 0.13$)

$$0.15\delta^2 \leq (\mathbf{w}_i - \mathbf{w}_j)^T M (\mathbf{w}_i - \mathbf{w}_j) \leq 4\delta^2$$

and furthermore every vector in this set also satisfies

$$\mathbf{1}^T\mathbf{w}_i = 0 \, .$$

Given this packing, we have

$$\max_{i,j} D_{\mathrm{KL}}(P_{\mathbf{w}_i}(\mathbf{y}) \| P_{\mathbf{w}_j}(\mathbf{y})) \leq \frac{4\delta^2}{\sigma^2}$$

and

$$\min_{i,j} (\mathbf{w}_i - \mathbf{w}_j)^T L (\mathbf{w}_i - \mathbf{w}_j) \geq 0.15\delta^2 \, .$$

Using Fano's inequality, we get

$$(\hat{\mathbf{w}} - \mathbf{w}^*)^T L (\hat{\mathbf{w}} - \mathbf{w}^*) \quad \geq \quad \frac{0.15}{2}\delta^2 \left(1 - \frac{\frac{4\delta^2}{\sigma^2} + \log 2}{0.13d}\right)$$

Choosing

$$\delta^2 = 0.0076\sigma^2 d \, ,$$

bounding $\frac{\log 2}{d} < 0.07$ whenever $d > 9$, and noting that

$$\mathfrak{M}_n^{\hat{\Sigma}}(\text{THURSTONE}) = \frac{1}{n}(\hat{\mathbf{w}} - \mathbf{w}^*)^T L (\hat{\mathbf{w}} - \mathbf{w}^*) \, ,$$

we get the desired result.

Upper Bounds: Define function ℓ as

$$\ell(\mathbf{w}) = \sum_{i=1}^{n} (y_i - \mathbf{x}_i^T\mathbf{w})^2 \, .$$

Consider the maximum likelihood estimator

$$\hat{\mathbf{w}} \in \arg\min_{\mathbf{w}} \ell(\mathbf{w}) \, .$$

The solution is not unique (since the objective is invariant to shifting of $\mathbf{w}$), and hence we impose an additional constraint

$$\hat{\mathbf{w}} \in \arg\min_{\mathbf{w}:\mathbf{w}^T\mathbf{1}=0} \ell(\mathbf{w}) \, .$$

This is a loss function for the maximum likelihood estimator (and needs to be minimized). Now, the gradient and Hessian of this loss function is

$$\nabla \ell(\mathbf{w}) = -2\sum_{i=1}^{n}(y_i - \mathbf{x}_i^T w)\mathbf{x}_i = -2X^T\boldsymbol{\epsilon}$$

$$\nabla^2 \ell(\mathbf{w}) = 2L \, .$$

The third and higher order derivatives of ℓ are zero.

Defining $\Delta := \hat{\mathbf{w}} - \mathbf{w}^*$, we have

$$\ell(\mathbf{w}^* + \boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle \nabla \ell(\mathbf{w}^*),\, \boldsymbol{\Delta} \rangle \quad = \quad 2\boldsymbol{\Delta}^T L \boldsymbol{\Delta} \, .$$

Also, since $\hat{\mathbf{w}}$ minimizes this loss function, we have

$$\begin{aligned} \ell(\mathbf{w}^* + \boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle \nabla \ell(\mathbf{w}^*),\, \boldsymbol{\Delta} \rangle &\leq -\langle \nabla \ell(\mathbf{w}^*),\, \boldsymbol{\Delta} \rangle \\ &\leq \sqrt{\nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*)} \sqrt{\boldsymbol{\Delta}^T L \boldsymbol{\Delta}} \end{aligned}$$

where the last equation follows from Lemma 7 proved below.

We shall now upper bound the quantity $\nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*)$. We have

$$\begin{aligned} \nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*) &= \boldsymbol{\epsilon}^T X K K X^T \boldsymbol{\epsilon} \\ &= \|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon}\|_2^2 \, . \end{aligned}$$

Now, $\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon} \sim N(0, \tilde{\Lambda}^{\frac{1}{2}} U L U^T \tilde{\Lambda}^{\frac{1}{2}})$ and hence

$$\begin{aligned} E[\|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon}\|_2^2] &= \mathrm{tr}(\tilde{\Lambda}^{\frac{1}{2}} U^T L U \tilde{\Lambda}^{\frac{1}{2}}) \\ &= d - 1 \, . \end{aligned}$$

We will use [1, Proposition 1] which says that for $\boldsymbol{\epsilon} \sim N(0, \sigma^2 I)$ and any matrix A,

$$P(\|A\boldsymbol{\epsilon}\|_2^2/\sigma^2 > \mathrm{tr}(A^T A) + 2\sqrt{\mathrm{tr}((A^T A)^2)t} + 2\|A^T A\|_2 t) \leq e^{-t} \qquad \forall\, t \geq 0 \tag{1}$$

In our setting, we have $A = \tilde{\Lambda}^{\frac{1}{2}} U^T X^T$ and hence

$$\mathrm{tr}(A^T A) = \mathrm{tr}(XKX^T) = \mathrm{tr}(LK) = d - 1 \, ,$$

$$\mathrm{tr}((A^T A)^2) = \mathrm{tr}(XKX^T XKX^T) = d - 1 \, ,$$

$$\|A\|_{\mathrm{op}}^2 = \arg\max_{\|\mathbf{v}\|_2=1} \mathbf{v}^T X K K X^T \mathbf{v} = \arg\max_{\|\mathbf{v}\|_2=1} \mathbf{v}^T V \Sigma^T U^T U \tilde{\Lambda} U^T U \Sigma V \mathbf{v} = \arg\max_{\|\mathbf{v}\|_2=1} \mathbf{v}^T \Sigma^T \tilde{\Lambda} \Sigma \mathbf{v} = 1$$

where the last equation follows from setting the singular value decomposition of X as $U\Sigma V^T$ and noting that by definition of K, we have $\Sigma^T \tilde{\Lambda} \Sigma = \sum_{i=1}^{d-1} \mathbf{e}_i \mathbf{e}_i^T$. Substituting these values, we have

$$P(\|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon}\|_2^2/\sigma^2 > (d-1) + 2\sqrt{(d-1)t} + 2t) \leq e^{-t} \qquad \forall\, t \geq 0$$

$$\Rightarrow P(\|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon}\|_2^2 > 2t\sigma^2 d) \leq e^{-t} \qquad \forall\, t \geq 1 \tag{2}$$

Putting everything together, we get

$$\sqrt{\Delta^T L \Delta} \leq \sqrt{\frac{d\sigma^2 t}{2}} \quad w.p. \geq 1 - e^{-t} \qquad \forall\, t \geq 1 \, .$$

Squaring both sides and substituting $\hat{\Sigma} = \frac{1}{n}L$, we get In terms of the standardized Laplacian, we have

$$\Delta^T \hat{\Sigma} \Delta \leq \frac{d\sigma^2 t}{2n} \quad w.p. \geq 1 - e^{-t} \qquad \forall\, t \geq 1\,.$$

Finally,

$$\frac{2n}{d\sigma^2}\mathbb{E}\left[\Delta^T \hat{\Sigma}\Delta\right] = \int_{t=0}^{\infty} \mathbb{P}\left(\frac{2n}{d\sigma^2}\Delta^T\hat{\Sigma}\Delta > t\right) \leq 1 + \int_{t=1}^{\infty} e^{-t} = 1 + \frac{1}{e}\,.$$

■

Proof of Theorem 2.B (THURSTONE**):**
Lower Bounds: Let Φ denote the c.d.f. of the standard Gaussian distribution and let ϕ denote its p.d.f. For any $\mathbf{w}_1$ and $\mathbf{w}_2$, the KL divergence between the distributions of $\mathbf{y}$ under $\mathbf{w}_1$ and $\mathbf{w}_2$ as the true values is

$$\begin{aligned}
D_{\mathrm{KL}}(P_{\mathbf{w}_1}(\mathbf{y})\|P_{\mathbf{w}_2}(\mathbf{y})) &= \sum_{i=1}^{n} \Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)\log\frac{\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)} + (1-\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma))\log\frac{1-\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)}{1-\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)} \\
&\leq \sum_{i=1}^{n} (\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma) - \Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma))\frac{\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)} \\
&\quad +((1-\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)) - (1-\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)))\frac{1-\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma)}{1-\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)} \\
&= \sum_{i=1}^{n} \frac{(\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma) - \Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma))^2}{\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma)(1-\Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma))} \\
&\leq \sum_{i=1}^{n} \frac{(\Phi(\mathbf{w}_1^T\mathbf{x}_i/\sigma) - \Phi(\mathbf{w}_2^T\mathbf{x}_i/\sigma))^2}{\Phi(2B/\sigma)(1-\Phi(2B/\sigma))} \\
&\leq \sum_{i=1}^{n} \frac{f(0)^2}{\Phi(2B/\sigma)(1-\Phi(2B/\sigma))}(\mathbf{w}_1^T\mathbf{x}_i/\sigma - \mathbf{w}_2^T\mathbf{x}_i/\sigma)^2\,. \\
&= \frac{1}{2\pi\sigma^2\Phi(2B/\sigma)(1-\Phi(2B/\sigma))}(\mathbf{w}_1-\mathbf{w}_2)^T L(\mathbf{w}_1-\mathbf{w}_2)
\end{aligned}$$

For any $\delta > 0$, Lemma 6 constructs a packing $\{\mathbf{w}_1, \ldots, \mathbf{w}_{e^{\beta d}}\}$ such that every pair of distinct vectors $\mathbf{w}_i$ and $\mathbf{w}_j$ in this packing satisfies (with $\alpha = 0.15$ and $\beta = 0.13$)

$$0.15\delta^2 \leq (\mathbf{w}_i - \mathbf{w}_j)^T M(\mathbf{w}_i - \mathbf{w}_j) \leq 4\delta^2$$

and furthermore every vector $\mathbf{w}_i$ in this set also satisfies

$$\mathbf{1}^T\mathbf{w}_i = 0\,.$$

Given this packing, we have

$$\max_{i,j} D_{\mathrm{KL}}(P_{\mathbf{w}_i}(\mathbf{y})\|P_{\mathbf{w}_j}(\mathbf{y})) \leq \frac{4\delta^2}{2\pi\Phi(2B/\sigma)(1-\Phi(2B/\sigma))\sigma^2}$$

and

$$\min_{i,j}(\mathbf{w}_i - \mathbf{w}_j)^T L(\mathbf{w}_i - \mathbf{w}_j) \geq 0.15\delta^2\,.$$

Using Fano's inequality, we get

$$(\hat{\mathbf{w}} - \mathbf{w}^*)^T L(\hat{\mathbf{w}} - \mathbf{w}^*) \quge \frac{0.15}{2}\delta^2\left(1 - \frac{\frac{4\delta^2}{2\pi\Phi(2B/\sigma)(1-\Phi(2B/\sigma))\sigma^2} + \log 2}{0.13d}\right)$$

Choosing

$$\delta^2 = 0.0076\sigma^2 d \times 2\pi\Phi(2B/\sigma)(1-\Phi(2B/\sigma))$$

bounding $\frac{\log 2}{d} < 0.07$ whenever $d > 9$, and noting that

$$\mathfrak{M}_n^{\hat{\Sigma}}(\textsc{Thurstone}) = \frac{1}{n}(\hat{\mathbf{w}} - \mathbf{w}^*)^T L(\hat{\mathbf{w}} - \mathbf{w}^*) \,,$$

we get the desired result. The only issue remaining to consider is the bounded assumption of $\mathbf{w}$, and this is verified below.

$$\begin{aligned}
\|\mathbf{w}\|_\infty &= \frac{\delta}{\sqrt{d}} \|U\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}\|_\infty \\
&\leq \frac{\delta}{\sqrt{d}} \sup_{\mathbf{u}:\|\mathbf{u}\|_2=1} \mathbf{u}^T \tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)} \\
&= \frac{\delta}{\sqrt{d}} \frac{(\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)})^T \tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}}{\|\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}\|_2} \\
&\leq \frac{\delta}{\sqrt{d}} \sqrt{\mathrm{tr}(\tilde{\Lambda})} \qquad (3) \\
&= \frac{\delta}{\sqrt{d}} \sqrt{\mathrm{tr}(K)} \\
&= \sqrt{0.00555\sigma^2 \times 2\pi\Phi(2B/\sigma)(1-\Phi(2B/\sigma))\frac{\mathrm{tr}(\hat{\Sigma}^\dagger)}{n}} \\
&\leq B \,,
\end{aligned}$$

where (3) follows from the fact that $\mathbf{w}^{(2)} \in \{-1, 0, 1\}^d$ and the final equation follows from our assumption relating n and $\mathrm{tr}(\hat{\Sigma}^\dagger)$.

Upper Bounds: Define function ℓ as

$$\ell(\mathbf{w}) = -\sum_{i=1}^n \left[\mathbf{1}\{y_i = 1\} \log \Phi(\mathbf{w}^T\mathbf{x}_i/\sigma) + \mathbf{1}\{y_i = -1\} \log(1 - \Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))\right] \,.$$

Consider the maximum likelihood estimator

$$\hat{\mathbf{w}} \in \operatorname*{arg\,min}_{\mathbf{w}:\mathbf{1}^T\mathbf{w}=0, \|\mathbf{w}\|_\infty \leq B} \ell(\mathbf{w}) \,.$$

The gradient and Hessian of this loss function are

$$\nabla\ell(\mathbf{w}) = \frac{-1}{\sigma}\sum_{i=1}^n \left[\mathbf{1}\{y_i = 1\}\frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)} - \mathbf{1}\{y_i = -1\}\frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}{1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}\right]\mathbf{x}_i \,,$$

and

$$\begin{aligned}
\nabla^2\ell(\mathbf{w}) &= \frac{1}{\sigma^2}\sum_{i=1}^n \left[\mathbf{1}\{y_i = 1\}\frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2 - \Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)\phi'(\mathbf{w}^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2}\right. \\
&+ \left. \mathbf{1}\{y_i = -1\}\frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2 + (1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))\phi'(\mathbf{w}^T\mathbf{x}_i/\sigma)}{(1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))^2}\right]\mathbf{x}_i\mathbf{x}_i^T \qquad (4)
\end{aligned}$$

respectively. The scalar in the summation is always non-negative (since Φ is log-concave), and hence maximum likelihood inference is a convex optimization problem.

Define

$$\begin{aligned}
c_1(\sigma, B) &:= \inf_{\mathbf{w}:\mathbf{1}^T\mathbf{w}=0, \|\mathbf{w}\|_\infty \le B, i\in[n]} \min\left\{ \frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2 - \Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)\phi'(\mathbf{w}^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2}, \right. \\
&\qquad \left. \frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2 + (1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))\phi'(\mathbf{w}^T\mathbf{x}_i/\sigma)}{(1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))^2} \right\} \\
&= \inf_{\mathbf{w}:\mathbf{1}^T\mathbf{w}=0, \|\mathbf{w}\|_\infty \le B, i\in[n]} \frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)^2 + (1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))\phi'(\mathbf{w}^T\mathbf{x}_i/\sigma)}{(1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))^2} \\
&= \inf_{t\in[-2B/\sigma, 2B/\sigma]} \frac{\phi(t)^2 + (1-\Phi(t))\phi'(t)}{(1-\Phi(t))^2} \\
&= \inf_{t\in[-2B/\sigma, 2B/\sigma]} \left(\frac{\phi(t)}{1-\Phi(t)}\right)^2 - t\frac{\phi(t)}{1-\Phi(t)} \\
&\ge \left(\frac{t+\sqrt{t^2+\frac{8}{\pi}}}{2}\right)^2 - t\left(\frac{t+\sqrt{t^2+4}}{2}\right) \\
&= \frac{2}{\pi} - \frac{t}{2}\left(\sqrt{t^2+4} - \sqrt{t^2+\frac{8}{\pi}}\right) \\
&= \frac{2}{\pi} - \frac{t}{2}\frac{(t^2+4)-(t^2+\frac{8}{\pi})}{\sqrt{t^2+4}+\sqrt{t^2+\frac{8}{\pi}}} \\
&= \frac{2}{\pi} - \left(2-\frac{4}{\pi}\right)\frac{t}{\sqrt{t^2+4}+\sqrt{t^2+\frac{8}{\pi}}} \\
&\ge \frac{4}{\pi} - 1\,.
\end{aligned}$$

Then for all $\mathbf{w}$ in the allowed set and any vector $\mathbf{v} \in \mathbb{R}^d$, we have

$$\mathbf{v}^T\nabla^2\ell(\mathbf{w})\mathbf{v} \;\ge\; \frac{c_1(\sigma,B)}{\sigma^2}\sum_{i=1}^n \mathbf{v}^T\mathbf{x}_i\mathbf{x}_i^T\,\mathbf{v}.$$

Defining $\Delta := \hat{\mathbf{w}} - \mathbf{w}^*$, we have

$$\begin{aligned}
\ell(\mathbf{w}^*+\boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle\nabla\ell(\mathbf{w}^*),\,\boldsymbol{\Delta}\rangle &\ge \boldsymbol{\Delta}^T\left(\frac{c_1(\sigma,B)}{\sigma^2}\sum_{i=1}^n \mathbf{x}_i\mathbf{x}_i^T\,.\right)\boldsymbol{\Delta} \\
&= \frac{c_1(\sigma,B)}{\sigma^2}\boldsymbol{\Delta}^T L\boldsymbol{\Delta}\,.
\end{aligned}$$

Also, since $\hat{\mathbf{w}}$ minimizes this loss function, we have

$$\begin{aligned}
\ell(\mathbf{w}^*+\boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle\nabla\ell(\mathbf{w}^*),\,\boldsymbol{\Delta}\rangle &\le -\langle\nabla\ell(\mathbf{w}^*),\,\boldsymbol{\Delta}\rangle \\
&\le \sqrt{\nabla\ell(\mathbf{w}^*)^T K\nabla\ell(\mathbf{w}^*)}\sqrt{\boldsymbol{\Delta}^T L\boldsymbol{\Delta}}
\end{aligned}$$

where the last equation follows from Lemma 7.

We will now upper bound the quantity $\nabla\ell(\mathbf{w}^*)^T K\nabla\ell(\mathbf{w}^*)$. Define independent random variables $\{\theta_i\}_{i=1}^n$ as

$$\theta_i = \begin{cases} \frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)} & \text{w.p.} \quad \Phi(\mathbf{w}^T\mathbf{x}_i/\sigma) \\ \frac{-\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}{1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)} & \text{w.p.} \quad 1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma) \end{cases}$$

and let $\boldsymbol{\theta}^T = [\theta_1, \cdots, \theta_n]$. Then

$$\nabla\ell(\mathbf{w}) = \frac{-1}{\sigma}X^T\boldsymbol{\theta}\,,$$

and

$$\begin{aligned}\nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*) &= \frac{1}{\sigma^2}\boldsymbol{\theta}^T X K X^T \boldsymbol{\theta} \\ &= \frac{1}{\sigma^2}\|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\theta}\|_2^2 \,.\end{aligned}$$

We will now apply [1, Theorem 2.1] which says that any random vector $\boldsymbol{\epsilon}$ that is zero-mean and sub-gaussian with parameter σ, and any matrix A, must satisfy (1). We will now set $\boldsymbol{\theta}$ as $\boldsymbol{\epsilon}$ and $\tilde{\Lambda}^{\frac{1}{2}} U^T X^T$ as A in (1). To this end, we see that

$$\mathbf{E}\left[\boldsymbol{\theta}\right] = \mathbf{0}$$

and by virtue of each coordinate being being bounded, $\boldsymbol{\theta}$ is sub-gaussian with parameter at most $c_2(B,\sigma)$ where $c_2(B,\sigma)$ is defined as

$$\begin{aligned}c_2(B,\sigma) &= \sup_{\mathbf{w}:\mathbf{1}^T\mathbf{w}=0, \|\mathbf{w}\|_\infty \le B} \max_{i\in[n]} \frac{\phi(\mathbf{w}^T\mathbf{x}_i/\sigma)}{\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma)(1-\Phi(\mathbf{w}^T\mathbf{x}_i/\sigma))} \\ &\le \frac{1}{\sqrt{2\pi}\Phi(2B/\sigma)(1-\Phi(2B/\sigma))} \,.\end{aligned}$$

Substituting these in (1) and following the simplifications of (2), we get

$$P\left(\|\tilde{\Lambda}^{\frac{1}{2}} U^T X^T \boldsymbol{\epsilon}\|_2^2 > 2t c_2(B,\sigma)^2 d\right) \le e^{-t} \qquad \forall\, t \ge 1 \,.$$

Putting everything together, we have

$$\boldsymbol{\Delta}^T L \boldsymbol{\Delta} \le \frac{c_2(B,\sigma)^2}{c_1(B,\sigma)^2} 2d\sigma^2 t \quad w.p. \ge 1-e^{-t} \quad \forall t \ge 1 \,.$$

Substituting the bounds on c_1 and c_2, and substituting $\hat{\Sigma} = \frac{1}{n}L$, we get

$$\boldsymbol{\Delta}^T \hat{\Sigma} \boldsymbol{\Delta} \le \frac{3.66}{\left(\Phi(2B/\sigma)(1-\Phi(2B/\sigma))\right)^2} \frac{d\sigma^2 t}{n} \quad w.p. \ge 1-e^{-t} \quad \forall t \ge 1 \,.$$

Converting this to a bound on $\mathbb{E}\left[\boldsymbol{\Delta}^T \hat{\Sigma} \boldsymbol{\Delta}\right]$ as done in the final step of the proof of Theorem 2.A gives the desired result. ■

Proof of Theorem 2.C (BTL)**:** For any differencing vector $\mathbf{x}$, let $a(\mathbf{x})$ be the index of the '1' in $\mathbf{x}$ and let $b(\mathbf{x})$ be the index of the '-1' in $\mathbf{x}$. Now define a function $\Psi : \mathbb{R}^d \times \{-1,0,1\}^d \to \mathbb{R}$, where the second argument is always a differencing vector, as

$$\Psi(\mathbf{w},\mathbf{x}) = \log\left(\exp\left(\frac{w_{a(\mathbf{x})}}{\sigma}\right) + \exp\left(\frac{w_{b(\mathbf{x})}}{\sigma}\right)\right) - \frac{w_{a(\mathbf{x})} + w_{b(\mathbf{x})}}{2\sigma} \,.$$

First, consider a single sample with observation y and differencing vector $\mathbf{x}$. We can rewrite the likelihood function of the BTL model as

$$P(y|w) = \exp\left(\frac{y}{2\sigma}(\mathbf{w})^T\mathbf{x} - \Psi(\mathbf{w},\mathbf{x})\right) .$$

Using this form, one can compute that

$$\begin{aligned}D_{KL}(P_{\mathbf{w}_1}(y)||P_{\mathbf{w}_2}(y)) = \frac{1}{2\sigma}\frac{1-e^{\frac{(\mathbf{w}_1)^T\mathbf{x}}{\sigma}}}{1+e^{\frac{(\mathbf{w}_1)^T\mathbf{x}}{\sigma}}}(\mathbf{w}_1-\mathbf{w}_2)^T\mathbf{x} \\ - \left((\mathbf{w}_1-\mathbf{w}_2)^T\nabla\Psi(\mathbf{w_1},\mathbf{x}) + (\mathbf{w}_1-\mathbf{w}_2)^T\nabla^2\Psi(\mathbf{w_3},\mathbf{x})(\mathbf{w}_1-\mathbf{w}_2)\right) \,, \quad (5)\end{aligned}$$

for some $\mathbf{w}_3$. One can evaluate that

$$\nabla\Psi(\mathbf{w_1},\mathbf{x}) = \frac{1}{2\sigma}\frac{1-e^{\frac{(\mathbf{w}_1)^T\mathbf{x}}{\sigma}}}{1+e^{\frac{(\mathbf{w}_1)^T\mathbf{x}}{\sigma}}}\mathbf{x}$$

and that

$$\begin{aligned}\nabla^2\Psi(\mathbf{w_3},\mathbf{x}) &= \frac{1}{2\sigma^2}\frac{1}{e^{\frac{(\mathbf{w}_3)^T\mathbf{x}}{\sigma}}+e^{\frac{-(\mathbf{w}_3)^T\mathbf{x}}{\sigma}}+2}\mathbf{x}\mathbf{x}^T \\ &\leq \frac{1}{8\sigma^2}\mathbf{x}\mathbf{x}^T\ .\end{aligned}$$

It follows that

$$D_{KL}(P_{\mathbf{w}_1}(y)||P_{\mathbf{w}_2}(y)) \leq \frac{1}{8\sigma^2}(\mathbf{w}_1-\mathbf{w}_2)^T\mathbf{x}\mathbf{x}^T(\mathbf{w}_1-\mathbf{w}_2)\ .$$

Aggregating this over all samples, and observing that the distribution of the observation is independent across samples, we get

$$D_{KL}(P_{\mathbf{w}_1}(y)||P_{\mathbf{w}_2}(y)) \leq \frac{1}{8\sigma^2}(\mathbf{w}_1-\mathbf{w}_2)^T L(\mathbf{w}_1-\mathbf{w}_2)\ .$$

For any $\delta > 0$, Lemma 6 constructs a packing $\{\mathbf{w}_1,\ldots,\mathbf{w}_{e^{\beta d}}\}$ such that every pair of distinct vectors $\mathbf{w}_i$ and $\mathbf{w}_j$ in this packing satisfies (with $\alpha = 0.15$ and $\beta = 0.13$)

$$0.15\delta^2 \leq (\mathbf{w}_i-\mathbf{w}_j)^T M(\mathbf{w}_i-\mathbf{w}_j) \leq 4\delta^2$$

and furthermore every vector in this set also satisfies

$$\mathbf{1}^T\mathbf{w}_i = 0\ .$$

Given this packing, we have

$$\max_{i,j} D_{KL}(P_{\mathbf{w}_1}(\mathbf{y})||P_{\mathbf{w}_2}(\mathbf{y})) \leq \frac{\delta^2}{2\sigma^2}$$

and

$$\min_{i,j}(\mathbf{w}_1-\mathbf{w}_2)^T L(\mathbf{w}_1-\mathbf{w}_2) \geq 0.15\delta^2\ . \tag{6}$$

Using Fano's inequality, we get

$$(\hat{\mathbf{w}}-\mathbf{w}^*)^T L(\hat{\mathbf{w}}-\mathbf{w}^*) \geq \frac{0.15}{2}\delta^2\left(1-\frac{\frac{\delta^2}{2\sigma^2}+\log 2}{0.13d}\right) \tag{7}$$

Choosing

$$\delta^2 = 0.06\sigma^2 d\ , \tag{8}$$

bounding $\frac{\log 2}{d} < 0.07$ whenever $d > 9$, and noting that

$$\mathfrak{M}_n^{\hat{\Sigma}}(\textsc{Thurstone}) = \frac{1}{n}(\hat{\mathbf{w}}-\mathbf{w}^*)^T L(\hat{\mathbf{w}}-\mathbf{w}^*)\ ,$$

we get the desired result. The only issue remaining to consider is the boundedness of $\mathbf{w}$, and this is verified below.

$$\begin{aligned}\|\mathbf{w}\|_\infty &= \frac{\delta}{\sqrt{d}}\|U\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}\|_\infty \\ &\leq \frac{\delta}{\sqrt{d}}\sup_{\mathbf{u}:\|\mathbf{u}\|_2=1}\mathbf{u}^T\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)} \\ &= \frac{\delta}{\sqrt{d}}\frac{(\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)})^T\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}}{\|\tilde{\Lambda}^{\frac{1}{2}}\mathbf{w}^{(2)}\|_2} \\ &\leq \frac{\delta}{\sqrt{d}}\sqrt{\mathrm{tr}(\tilde{\Lambda})} \\ &= \frac{\delta}{\sqrt{d}}\sqrt{\mathrm{tr}(K)} \\ &= \sqrt{0.04467\sigma^2\frac{\mathrm{tr}(\hat{\Sigma}^\dagger)}{n}} \\ &\leq B\ ,\end{aligned} \tag{9}$$

where (9) follows from the fact that $\mathbf{w}^{(2)} \in \{-1, 0, 1\}^d$ and the final equation follows from our assumption relating n and $\mathrm{tr}(\hat{\Sigma}^\dagger)$.

Upper Bounds: Define function ℓ as

$$\ell(\mathbf{w}) = \sum_{i=1}^{n} \log\left(1 + \exp\left(\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}\right)\right) .$$

Consider the maximum likelihood estimator

$$\hat{\mathbf{w}} \in \underset{\mathbf{w}:\mathbf{1}^T\mathbf{w}=0, \|\mathbf{w}\|_\infty \le B}{\arg\min} \ell(\mathbf{w}) .$$

The gradient and Hessian of this loss function are

$$\nabla \ell(\mathbf{w}) = \frac{1}{\sigma} \sum_{i=1}^{n} \frac{-y_i e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}}{1 + e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}} \mathbf{x}_i$$

$$\nabla^2 \ell(\mathbf{w}) \quad = \quad \frac{1}{\sigma^2} \sum_{i=1}^{n} \frac{e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}}{\left(1 + e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}\right)^2} \mathbf{x}_i \mathbf{x}_i^T .$$

One can see that the Hessian is positive semi-definite, making function ℓ a convex function.

Then for all $\mathbf{w}$ in the allowed set, any observation $y_i \in \{-1, 1\}$ and any differencing vector $\mathbf{x}_i$, it must be that

$$\frac{e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}}{\left(1 + e^{\frac{-y_i \mathbf{w}^T \mathbf{x}_i}{\sigma}}\right)^2} \ge \frac{1}{\left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^2} .$$

Defining $\Delta := \hat{\mathbf{w}} - \mathbf{w}^*$, we have

$$\begin{aligned} \ell(\mathbf{w}^* + \boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle \nabla \ell(\mathbf{w}^*), \boldsymbol{\Delta} \rangle &\ge \boldsymbol{\Delta}^T \nabla^2 \ell(\mathbf{w}) \boldsymbol{\Delta} \\ &\ge \boldsymbol{\Delta}^T \left(\frac{1}{\sigma^2 \left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^2} \sum_{i=1}^{n} \mathbf{x}_i \mathbf{x}_i^T . \right) \boldsymbol{\Delta} \\ &= \frac{1}{\sigma^2 \left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^2} \boldsymbol{\Delta}^T L \boldsymbol{\Delta} . \end{aligned}$$

Also, since $\hat{\mathbf{w}}$ minimizes this loss function, we have

$$\begin{aligned} \ell(\mathbf{w}^* + \boldsymbol{\Delta}) - \ell(\mathbf{w}^*) - \langle \nabla \ell(\mathbf{w}^*), \boldsymbol{\Delta} \rangle &\le -\langle \nabla \ell(\mathbf{w}^*), \boldsymbol{\Delta} \rangle \\ &\le \sqrt{\nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*)} \sqrt{\boldsymbol{\Delta}^T L \boldsymbol{\Delta}} \end{aligned}$$

where the last equation follows from Lemma 7.

We will now upper bound the quantity $\nabla \ell(\mathbf{w}^*)^T K \nabla \ell(\mathbf{w}^*)$. Define independent random variables $\{\theta_i\}_{i=1}^n$ as

$$\theta_i = \begin{cases} \frac{-e^{\frac{\mathbf{w}^T \mathbf{x}_i}{\sigma}}}{1+e^{\frac{\mathbf{w}^T \mathbf{x}_i}{\sigma}}} & \text{w.p.} \quad \frac{1}{1+e^{\frac{\mathbf{w}^T \mathbf{x}_i}{\sigma}}} \\ \frac{e^{\frac{-\mathbf{w}^T \mathbf{x}_i}{\sigma}}}{1+e^{\frac{-\mathbf{w}^T \mathbf{x}_i}{\sigma}}} & \text{w.p.} \quad \frac{1}{1+e^{\frac{-\mathbf{w}^T \mathbf{x}_i}{\sigma}}} \end{cases}$$

and let $\boldsymbol{\theta}^T = [\theta_1, \cdots, \theta_n]$. Then

$$\nabla \ell(\mathbf{w}) = \frac{-1}{\sigma} X^T \boldsymbol{\theta} ,$$

and

$$\begin{aligned}\nabla\ell(\mathbf{w}^*)^T K \nabla\ell(\mathbf{w}^*) &= \frac{1}{\sigma^2}\boldsymbol{\theta}^T XKX^T\boldsymbol{\theta} \\ &= \frac{1}{\sigma^2}\|\tilde{\Lambda}^{\frac{1}{2}}U^T X^T\boldsymbol{\theta}\|_2^2 \,.\end{aligned} \tag{10}$$

We will now apply [1, Theorem 2.1] which says that any random vector $\boldsymbol{\epsilon}$ that is zero-mean and sub-gaussian with parameter σ, and any matrix A, must satisfy (1). We will now set $\boldsymbol{\theta}$ as $\boldsymbol{\epsilon}$ and $\tilde{\Lambda}^{\frac{1}{2}}U^T X^T$ as A in (1). To this end, we see that

$$\mathbf{E}\left[\boldsymbol{\theta}\right] = \mathbf{0}$$

and by virtue of each coordinate being being bounded, $\boldsymbol{\theta}$ is sub-gaussian. The sub-gaussianity parameter is upper bounded by

$$\frac{e^{\frac{-\mathbf{w}^T\mathbf{x}_i}{\sigma}}}{1+e^{\frac{-\mathbf{w}^T\mathbf{x}_i}{\sigma}}} - \frac{-e^{\frac{\mathbf{w}^T\mathbf{x}_i}{\sigma}}}{1+e^{\frac{\mathbf{w}^T\mathbf{x}_i}{\sigma}}} = 1 \,.$$

Substituting these in (1) and following the simplifications of (2), we get

$$P(\|\tilde{\Lambda}^{\frac{1}{2}}U^T X^T\boldsymbol{\epsilon}\|_2^2 > 2td) \leq e^{-t} \qquad \forall\, t \geq 1 \,.$$

Putting everything together, we have

$$\sqrt{\boldsymbol{\Delta}^T L\boldsymbol{\Delta}} \leq \left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^2 \sqrt{2d\sigma^2 t} \quad w.p. \geq 1 - e^{-t} \quad \forall t \geq 1 \,.$$

Squaring and substituting $\hat{\Sigma} = \frac{1}{n}L$, we get

$$\boldsymbol{\Delta}^T\hat{\Sigma}\boldsymbol{\Delta} \leq \left(e^{\frac{B}{\sigma}} + e^{\frac{-B}{\sigma}}\right)^4 \frac{d\sigma^2 t}{n} \quad w.p. \geq 1 - e^{-t} \quad \forall t \geq 1 \,.$$

Converting this to a bound on $\mathbb{E}\left[\boldsymbol{\Delta}^T\hat{\Sigma}\boldsymbol{\Delta}\right]$ as done in the final step of the proof of Theorem 2.A gives the desired result. ■

Proof of Lemma 6: First construct a set of $e^{\beta d}$ vectors $\{\mathbf{w}_1^{(1)}, \ldots, \mathbf{w}_{e^{\beta d}}^{(1)}\}$, each belonging to $\{-1, +1\}^{d-1}$ such that

$$\min_{i\neq j} \text{Hamming-distance}(\mathbf{w}_i^{(1)} - \mathbf{w}_j^{(1)}) \geq \alpha d \,.$$

The existence of such a set is guaranteed by the Gilbert-Varshamov bound, which guarantees existence of a (binary) code of length $(d-1)$, minimum Hamming distance αd, and the number of code words at least

$$\begin{aligned}\frac{2^{d-1}}{\sum_{\ell=0}^{\alpha d-1}\binom{d-1}{\ell}} &\geq 2^{d-1}\left(\frac{\alpha d-1}{(d-1)e}\right)^{\alpha d-1} \\ &= e^{(d-1)\log 2 + (\alpha d-1)\log\left(\frac{\alpha d-1}{ed-e}\right)} \\ &\geq e^{\frac{d}{2}\left(\log 2 + \alpha\log\left(\frac{\alpha}{e}\right)\right)} \\ &= e^{\beta d} \,.\end{aligned}$$

It follows from the construction that for every pair of distinct vectors in this set,

$$4\alpha d \leq (\mathbf{w}_i^{(1)} - \mathbf{w}_j^{(1)})^T(\mathbf{w}_i^{(1)} - \mathbf{w}_j^{(1)}) \leq 4d \,.$$

Now construct a second set of $e^{\beta d}$ vectors $\{\mathbf{w}_1^{(2)}, \ldots, \mathbf{w}_{e^{\beta d}}^{(2)}\}$, each of length d, as

$$\left(\mathbf{w}_i^{(2)}\right)^T = \left[\left(\mathbf{w}_i^{(1)}\right)^T \quad 0\right]^T \qquad \forall i \,.$$

It is easy to see that every pair of distinct vectors in this set satisfies

$$4\alpha d \leq (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)})^T(\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)}) \leq 4d \; .$$

Finally, construct a third set of $e^{\beta d}$ vectors $\{\mathbf{w}_1, \ldots, \mathbf{w}_{e^{\beta d}}\}$, each of length d, as

$$\mathbf{w}_i = \frac{\delta}{\sqrt{d}} U \tilde{\Lambda}^{\frac{1}{2}} \mathbf{w}_i^{(2)} \quad \forall i \; .$$

For any vector in this set

$$\begin{aligned} \mathbf{1}^T \mathbf{w}_i &= \frac{\delta}{\sqrt{d}} \mathbf{1}^T U \tilde{\Lambda}^{\frac{1}{2}} \mathbf{w}_i^{(2)} \\ &= \frac{\delta}{\sqrt{d}} \mathbf{e}_d \mathbf{e}_d^T \tilde{\Lambda}^{\frac{1}{2}} \mathbf{w}_i^{(2)} \\ &= 0 \; . \end{aligned}$$

For any pair of vectors in this set,

$$\begin{aligned} (\mathbf{w}_i - \mathbf{w}_j)^T L (\mathbf{w}_i - \mathbf{w}_j) &= \frac{\delta^2}{d} (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)})^T \tilde{\Lambda}^{\frac{1}{2}} U^T L U \tilde{\Lambda}^{\frac{1}{2}} (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)}) \\ &= \frac{\delta^2}{d} (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)})^T \tilde{\Lambda}^{\frac{1}{2}} \Lambda \tilde{\Lambda}^{\frac{1}{2}} (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)}) \\ &= \frac{\delta^2}{d} (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)})^T (\mathbf{w}_i^{(2)} - \mathbf{w}_j^{(2)}) \end{aligned}$$

where the last step makes use of the fact that the last coordinate of each vector in the set $\{\mathbf{w}_1^{(2)}, \ldots, \mathbf{w}_{e^{\beta d}}^{(2)}\}$ is zero. It follows that

$$4\alpha\delta^2 \leq (\mathbf{w}_i - \mathbf{w}_j)^T L (\mathbf{w}_i - \mathbf{w}_j) \leq 4\delta^2 \; .$$

■

Proof of Lemma 7: Consider the singular value decompositions $L = U\Lambda U^T$, $K = U\tilde{\Lambda}U^T$. Let $\tilde{\mathbf{x}} := \Lambda^{\frac{1}{2}} U^T \mathbf{x}$ and $\tilde{\mathbf{y}} := \tilde{\Lambda}^{\frac{1}{2}} U^T \mathbf{y}$. Then

$$\begin{aligned} \sqrt{\mathbf{x}^T L \mathbf{x}} \sqrt{\mathbf{y}^T K \mathbf{y}} &= \sqrt{\mathbf{x}^T U \Lambda U^T \mathbf{x}} \sqrt{\mathbf{y}^T U \tilde{\Lambda} U^T \mathbf{y}} \\ &= \|\tilde{\mathbf{x}}\|_2 \|\tilde{\mathbf{y}}\|_2 \\ &\geq \tilde{\mathbf{x}}^T \tilde{\mathbf{y}} \\ &= \mathbf{x}^T U \Lambda^{\frac{1}{2}} \tilde{\Lambda}^{\frac{1}{2}} U^T \mathbf{y} \\ &= \mathbf{x}^T U U^T \mathbf{y} \quad \text{(since } \mathbf{x} \perp \text{nullspace}(L)\text{)} \\ &= \mathbf{x}^T \mathbf{y} \; . \end{aligned}$$

■

References

[1] D. Hsu, S. M. Kakade, and T. Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electron. Commun. Probab*, 17(52):6, 2012.

[2] E. Lehmann and G. Casella. *Theory of Point Estimation*. Springer Texts in Statistics. 1998.

[3] A. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. 2008.